12

Learning Who Your Friends Are: Bayesian Games and Private Information

The philosophers kick up the dust and then complain that they cannot see.

Spinoza

12.1 Private Information and Bayesian Games

Occasionally we have come across a game in which there is *private information*: some pieces of relevant information are known by some players but not by others *even before any player has made a move in the game.*[1] For instance in *The Truth Game* (§3.19) the car dealer knows the quality of the car but the prospective buyer does not. In *One-card, Two-round Poker with Bluffing* (§4.11), One knows whether the card is high or low, but Two does not. In *Real Men Don't Eat Quiche* (§4.39), Clem knows whether he is Tough or a Wimp, but the Dark Stranger does not.

Private information can take several forms. First, the payoffs to a particular player may depend on the player's type, as in the case of Clem who likes quiche only if he's a Wimp. Second, a player's type may affect the payoffs to other players, so for instance a strong opponent may inflict more harm than a weak opponent. Third, various states of the world may be known to some players but not others, as in the Truth Game, where only the dealer knows the quality of the car. Finally, some players may be precommited to making some choices, and precluded from making others. For instance, in *Predatory Pricing: Pooling and Separating Equilibria* (§12.10), a 'sane' firm may share the market with a new entrant if profitable to do so, while a 'crazy' incumbent always preys on a new entrant, even if this leads to a lower payoff.

[1] Such games are usually called games of *incomplete information* in the literature. But this seems misleading, since unless information is private in the sense of being *asymmetrically distributed* among players, the reasoning we are about to undertake is vacuous.

VISIT…

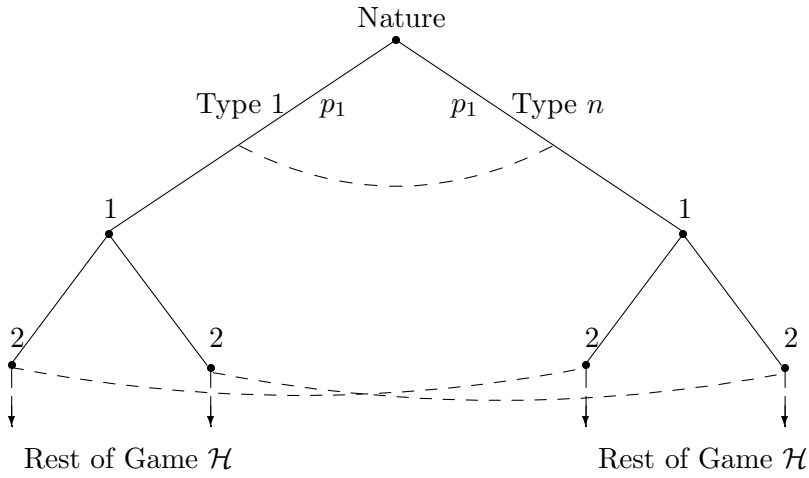

Figure 12.1: Modeling Games with Private Information: The Harsanyi Transformation

There is a straightforward way to model games with private information, due to Harsanyi (1967). Notice that in all three games mentioned above (§3.19, §4.11, §4.39), Nature moves first, determining the private information available to one of the players (high/low quality, high/low card, Wimp/Tough). Then a complete copy of a game tree $\mathcal{H}$ describing the player's moves is attached to the to the tail of each branch corresponding to Nature's moves, and the nodes from each copy of $\mathcal{H}$ where a player is ignorant of Nature's move are collected into a single information set for that player. In the general case we do exactly the same thing.

Formally describing the general case leads to a flurry of mathematical symbols that belies the simplicity of the procedure, so I will describe only the case where only Player One has private information. Let $\mathcal{H}$ be a game in which Player One may be one of a number of *types*, each 'type' representing the private information One possesses. We make a new game $\mathcal{G}$ as illustrated in Figure 12.1.

First, at the root node of $\mathcal{G}$, Nature chooses a type $\{1, \ldots, n\}$ for Player One, using a probability distribution $\{p_1, \ldots, p_n\}$ over types that is common knowledge to all the players. For each such choice—i.e., for each branch emanating from the root node—we append a copy of $\mathcal{H}$. Second, for each player i other than Nature and Player One, we note that each information set of player i in $\mathcal{H}$ has as many copies in $\mathcal{G}$ as there are Player One types. We collect all these information sets into one information set of $\mathcal{G}$, and assign this to player i. Thus Player One now knows his type, but other players only know the frequency distribution

of Player One types.[2] Finally, we define a strategy for Player One in $\mathcal{G}$ to be a choice of a strategy of $\mathcal{H}$ for each of the types Player One can have. The strategies for the other players in $\mathcal{G}$ are the same as in $\mathcal{H}$.

Games with private information often have many implausible Nash equilibria, and subgame perfection (as well as its mild generalization fuzzy subgame perfection—see §5.13), is rarely available to reject them because information sets are so pervasive that games with private information rarely have subgames. Moreover, there are often many mixed strategies that dictate exactly the same player choices, so the standard normal form notation obscures the game's behavioral regularities. For these reasons, in analyzing games with private information we often use the concept of a *perfect Nash equilibrium in behavioral strategies* in place of the Nash equilibrium in mixed strategies. This equilibrium concept is developed in §4.44 and §5.14, so if you have not yet studied those sections, now is the time to do so. For the impatient, however, a brief summary should suffice. We assume perfect recall (§1.2) throughout.

A *behavioral strategy* p_i for player i in an extensive form game $\mathcal{G}$ is a probability distribution p_ν over the player's action set at each of the player's information sets $\nu \in \mathcal{N}_i$, and a *behavioral strategy profile* p is the choice of a behavioral strategy for each player. Also included in p are the probability distributions over the nodes where Nature moves in $\mathcal{G}$. The *payoffs to a behavioral strategy profile* p are the expected payoffs to the players, assuming they use the probabilities given by p.

Starting from a mixed strategy σ of the normal form of $\mathcal{G}$, there is a natural way of deriving a *behavioral strategy representation* p of σ, such that the probabilities with which node are chosen from information sets are the same for p and σ, and p and σ have the same payoffs. Moreover, any behavioral strategy p has a *mixed strategy representation* σ that (a) induces the same probabilities on information sets, and (b) has the same payoffs as p. The association between mixed strategies and behavioral strategies is not one-to-one, however, because many mixed strategies can have the same behavioral representation.

We say a behavioral strategy profile $p = (p_1, \ldots, p_n)$ is a *Nash equilibrium in behavioral strategies* if for each player i, p_i is a best response to p_{-i}. This is equivalent to saying that the mixed strategy representation σ of p is a Nash equilibrium of the normal form of $\mathcal{G}$.

The behavioral strategy concept allows us to define a new, and often

[2]It may sound silly to say that Player One does not know such personal characteristics as sex and age until 'revealed' by Nature's choice, but this 'veil of ignorance' assumption is in fact purely formal.

appropriate, best response concept, call the *local best response*. Suppose i is choosing at node $a \in \nu$. The expected payoff to using p_ν given $p_{-\nu}$, is then

$$\pi_i(p_\nu|a, p_{-\nu}) = \sum_{t\in T} \mathrm{P}[a, t|p_\nu, p_{-\nu}]\pi_i(t).$$

where T is the set of terminal nodes, $\pi_i(t)$ is the payoff to i at terminal node $t \in T$, and $\mathrm{P}[a, t|p]$ is the probability of reaching t from a given p. If $\mathrm{P}[a|p_{-\nu}, \nu]$ is the probability of being at node a, given that i is choosing at ν and p is being played, then i should choose p_ν to maximize

$$\pi_i(p_\nu|p_{-\nu}, \nu) = \sum_{a\in\nu} \mathrm{P}[a|\nu, p_{-\nu}] \sum_{t\in T} \mathrm{P}[a, t|p_\nu, p_{-\nu}]\pi_i(t).$$

Moreover $\mathrm{P}[a|p, \nu]$ is given by *Bayesian updating*:

$$\mathrm{P}[a|\nu, p_{-\nu}] = \frac{\mathrm{P}[a|p_{-\nu}]}{\mathrm{P}[\nu|p_{-\nu}]}$$

where $\mathrm{P}[\nu|p_{-\nu}] = \sum_{a\in\nu} \mathrm{P}[a|p_{-\nu}]$. Notice that when $\mathrm{P}[\nu|p_{-\nu}] = 0$ (i.e., ν is not reached using $p_{-\nu}$) this definition does not make any sense, and we define $\mathrm{P}[a|\nu, p_{-\nu}]$ arbitrarily. We then have the following *Fundamental Theorem on Bayesian Updating in Extensive Form Games:*

Theorem 12.1. *If p is a Nash equilibrium of the extensive form game $\mathcal{G}$ and if ν is an information set that is reached with positive probability given p, then p_ν is a local best response to $p_{-\nu}$.*

This theorem gives us a quite convenient way of calculating Nash equilibria in behavioral strategies. In fact, in the literature this Nash equilibrium concept is called a *perfect Bayesian Nash equilibrium*, and I will do so as well henceforth.

If an information set ν is not reached under p, it is still plausible that the local best response condition should hold at ν for all, or for an "appropriate subset" of probability distributions $\mathrm{P}[a|\nu, p_{-\nu}]$ over the nodes at ν. The reasons for this are the same as for subgame perfection: should an agent in fact choose at ν, he may use a best response, so other agents will do well not to heed incredible threats that entail avoiding ν. How we determine "appropriate subsets" of probability distributions over unreachable information sets is, however, a gray area of game theory. One strong criterion is Selten's *trembling hand perfection* (§5.16), which implies that agents always behave as though there is a slight probability that an information set will be reached, and play best responses even to very highly improbable states of affairs.

The trembling hand criterion is however much too strong in many situations, because *evolutionary processes ignore low probability events.* It follows *a fortiori* that evolutionary processes ignore zero probability events. For a codfish, the probability of being impaled on a fisher's hook is such a low probability event that is simply not registered in the fish's perceptual or behavioral repertoire. It is folly to represent this situation as one in which the cod does not know which node (prey versus fishhook) it occupies in an information set. Reed warblers raise cuckoo eggs as their own not because they occupy an information set at which is very low probability that an egg is not their own, but because evolutionary pressures have not driven the species to make such distinctions. Nor is there any strong reason to believe that such pressures will develop, since cuckoos have evolved so as to change their hosts frequently. Human beings may be more susceptible to trembling hand perfection, but I would not bet my bank account on it. If you give people offers they cannot refuse, they often refuse anyway. If you put people in unfamiliar situations and depend on their reacting optimally, or even well, to these situations, you are not long for this world.

These considerations should not be taken as asserting that there are certain inherent limitations to game theory, but rather that we need new modeling concepts do deal with low probability events. For instance, perhaps we should assume that players partition the event world into sets whose probability is at least p_*, and choose best responses to based on these partitions. A mutant, however, can repartition the event space without violating the p_* minimum, to 'invade' a Nash equilibrium based on a given partition. Since the strategy space is no longer finite (or closed and bounded), there need not necessarily be a Nash equilibrium of the metagame where partitions evolve, so the description of strategic interaction becomes one of temporary equilibria with shifting ways of carving up reality.

12.2 The Role of Beliefs in Games with Private Information

If you have previously studied game theory, you will no doubt have noticed that our treatment of Bayesian updating in games with private information has not relied at all on the concept of 'beliefs.' This is not an oversight, but rather a natural side-effect of our evolutionary perspective. Classical game theory takes the decision processes of the *rational actor* as central, whereas evolutionary game theory takes the *behavior* of corporeal actors as central and models the game's evolution

using replicator dynamics, diffusion processes, and other population-level phenomena. Beliefs in such a framework are the *explicandum*, not the *explicans*—a shorthand way of expressing a behavioral regularity rather than the source of the regularity. There is absolutely no need to introduce the concept of beliefs into the general theory of games with private information.[3]

It may be comforting to the classically-minded to have a notion that agents have beliefs at the start of the game (called *priors*), which they update using Bayes rule in the course of the game, and to describe the process as *sequential rationality*, but the notion is completely unnecessary. The assumption that agents know the frequency distribution of types is no more or less in need of justification that the assumption that they know the rules of the game, the payoffs and the decision tree. It is an assumption that is justified by its ability to explain the behavior of agents in social life. Moreover, the 'belief' concept invites all sorts of philosophical nonsense of the type parodied in the Spinoza quote at the head of the chapter.

Consider, for example, the great debate over the status of the assumption that 'priors' are common knowledge. Harsanyi and other venerable game theorists (for instance Aumann, 1987) have defended this *common prior assumption* on the grounds that 'rational agents' faced with the same information must have the same priors. Others dispute this (e.g., Kreps 1990), and indeed there may be compelling technical arguments as to why, even in the case of costless information processing with no informational capacity constraints, identical information need not entail identical beliefs (Kurz 1997). Moreover, people appear to have consistent an predictable judgmental biases. There is, for instance, a considerable body of evidence that in bargaining and litigation people have "self-serving biases" that do not simply reflect informational heterogeneity (Loewenstein, Issacharoff, Camerer, and Babcock 1992, Babcock, Loewenstein, Issacharoff, and Camerer 1995, Babcock and Loewenstein 1997).

The common prior assumption has also been defended by arguing that when there are differences in priors among agents, we should expand the number of agents types to include such differences—Mertens and Zamir (1985) show that this is always possible—so common priors then hold

[3]There is one well-known equilibrium refinement concept that depends integrally upon a concept of beliefs, that of *sequential equilibrium* (Kreps and Wilson 1982), but the concept is virtually identical to (actually slightly weaker than) that of trembling hand perfection, which does not use a concept of beliefs.

in the larger game. However if we create new types to capture heterogeneous beliefs, the game becomes correspondingly more complex, and the assumption that the game structure and the frequency distribution of types are common knowledge becomes correspondingly less plausible. There may be cases where people do not have common priors over the parameters of the game, but they *do* have common priors concerning the frequency distribution of *differences* in priors, but probably not on this planet.

From the evolutionary perspective, there is no clear value in using game theory to analyze situations that have not been subjected to an extensive evolutionary process, after which it might be plausible that agents have *correct*, and hence *common* priors, on grounds that agents that have inaccurate maps of the world are less likely to survive than agents who have accurate maps. But then again, perhaps not.

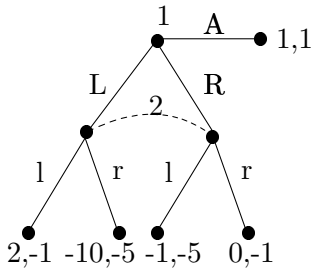

As an example of these ideas, consider the little game at the right, adapted from Kreps and Wilson (1982). There are two perfect Bayesian equilibria, (A, r) and (L, l), where the former requires weighting the probability of being at left-hand node of Player 2's information set at least 0.5. It might be thought that the former equilibrium is 'irrational,' since if Player 2 does indeed choose, then Player 1 must have played L, since by playing R, Player 1 cannot do as well as playing A, whatever Player 2 does. But this is to force more structure on the game than it in fact possesses. It might be argued that if Player 2 in fact chooses at all, then Player 1 must have played L, since playing R ensures that Player 1 does worse that when playing A. Therefore the only 'rational' move for player 2 is l. Knowing this, Player 1 should play L, and the (A, r) equilibrium cannot be played.

However consider Figure 12.2, which shows a simulation of the game with 500 randomly programmed agents of each type, who are randomly paired in each period, and 2% of which die and are replaced by more ' fit' agents (i.e., ones with higher-payoff strategies) after every 100 periods. With random mutations, the (A, r) equilibrium eventually emerges. This outcome is typical of the random mutation case. If we bias mutations in favor of L, then we do obtain the (L, l) equilibrium as exhibited in Figure 12.3. We could develop a model in which the mutation rate is itself subject to evolutionary forces, in which either equilibrium could predominate, depending on the precise structure of the larger model.

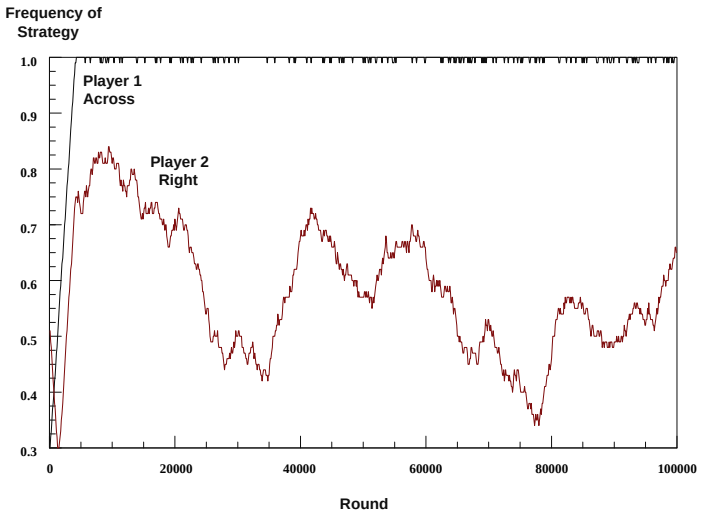

Figure 12.2: Simulation with Random Mutations

12.3 Haggling at the Bazaar

Consider a seller facing a buyer in a two-period game. In period 1 the seller makes an offer to sell at price p_1, and the buyer accepts or rejects. If the buyer accepts, the exchange is made and the game is over. Otherwise, the seller makes another offer p_2, and the buyer accepts or rejects. If he accepts in the second period, the exchange is made. Otherwise no trade occurs. The game tree is depicted in Figure 12.4.

Suppose the reservation price of the good to the seller is s, and the value to the buyer is b. Suppose buyer and seller have discount factors δ_s and δ_b for trades that are made in the second period. The value b to the buyer is unknown to the seller, but the seller believes that with probability π it is b_h, and with probability 1-π it is b_l, where $b_h > b_l > s$ (so the seller would gain from transacting with either type). Suppose that the parameters of the problem are such that if the seller did not get a second chance to make an offer, he would charge the lesser amount b_l, since otherwise the seller will ask the higher price in both periods, and

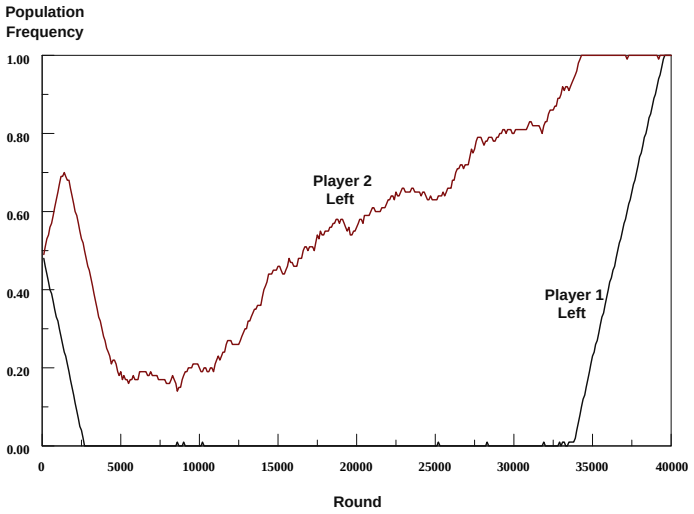

Figure 12.3: Simulation with Mutations Biased Towards L

the problem is uninteresting.

(a) Show that the payoffs are $[p_1, b - p_1]$, $[\delta_s p_2, \delta_b(b - p_2)]$, and $[\delta_s s, 0]$ in case the buyer accepts on the first round, the second round, and neither round, respectively. Draw the game tree and label the payoffs appropriately;

(b) Explain why the statement "if the seller did not get a second chance to make an offer, he would charge the lesser amount b_l" is equivalent to the statement that $\pi b_h + (1-\pi)s \leq b_l$, or simply $\pi \leq (b_l - s)/(b_h - s)$;

(c) Suppose $b_h \geq p_1 > b_l$, so there is some chance of getting to the second round. Let $\mu(p_1)$ be the seller's posterior probability that $b = b_h$, given that the buyer refused on round 1 at price p_1. Show that $\mu(p_1) = 1 - (1 - \pi)/(1 - \mathrm{P}(H))$ where H is the event "the buyer accepts price p_1 on round 1," and $\mathrm{P}(H)$ is the probability that H occurs. Show that this implies $\mu(p_1) \leq \pi$; i.e., if we get

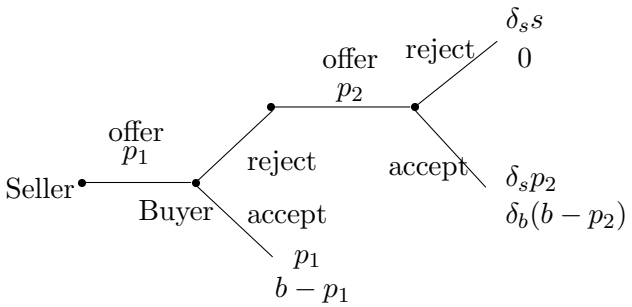

Figure 12.4: Haggling at the Bazaar

to the second round, the seller's posterior probability of the event $\{b = b_h\}$ cannot increase.

(d) Thus conclude that if we reach a second round, the seller will offer $p_2 = b_l$.

(e) Now roll back the game tree to the one-stage game in which the seller offers price p_1, the payoff for the buyer strategy "accept the offer" is $(p_1, b - p_1)$ and the payoff to the buyer strategy "reject the offer" is $(\delta_s b_l, \delta_b(b - b_l))$.

(f) Define $p^* = b_h - \delta_b(b_h - b_l)$. Show that the only undominated strategies for the seller are $p_1 = b_l$ and $p_1 = p^*$.

(g) Show that the seller chooses $p_1 = b_l$ or $p_1 = p^*$ according as b_l is greater than or less than $\pi p^* + (1 - \pi)\delta_s b_l$. Show that this reduces to the inequality

$$b_h \geq \left(1 + \frac{1 - \pi}{\pi} \frac{\delta_s}{1 - \delta_b}\right) b_l. \tag{12.1}$$

(h) Conclude that there is a unique perfect Bayesian equilibrium. What does equation (12.1) tell you about the conditions favoring the buyer or the seller?

(i) Suppose the parameters of the problem are such that if the seller did not get a second chance to make an offer, he would charge the *greater* amount b_h. Show that there is no Bayesian perfect Nash equilibrium in pure strategies.

12.4 Adverse Selection

In traditional models of the market for a good, supply and demand schedules intersect, determining the price. But in traditional models the characteristics of the good being traded are contractually specified or otherwise guaranteed fixed as the price is varied. The real world is often not so accommodating. Where the characteristics cannot be fixed, it is possible that some valued aspect of the good, which we may call its *quality*, will vary systematically with the price. We say there is *adverse selection* if the quality of the supply declines as the price declines. Here are some possible examples:

Labor: The wage in a traditional labor market is set to equate supply and demand. But this assumes the quality of workers is known by employers. Suppose not, and a firm wishing to hire an employee must incur costs to evaluate applicants. The firm advertises a wage w for the job and evaluates applicants until a successful applicant is found. If the firm increases the wage offer, the pool of qualified applicants could increase, because some qualified workers who would not apply at the lower wage do apply at the higher. Thus there is adverse selection in the labor market. The profit-maximizing firm will raise the wage offer to the point where the costs of paying a higher wage are just offset by the gains from lower transactions costs. Notice that the wage will be higher than the market-clearing wage in this case, so there will be an *excess supply of labor* in (Nash) equilibrium.

Capital: The interest rate on a loan in a traditional capital market is set to where the demand and supply for capital are equal. But this assumes that the quality of a borrower (the probability that the borrower will repay the loan) is contractually specifiable or otherwise known to a lender. Suppose a bank posts an interest rate r. If there is an excess demand for loans at that interest rate, the bank may be tempted to raise the interest rate. By so doing, some relatively low risk, low expected return projects would drop out of the demand, since they could not make a positive profit at a higher interest rate. But high risk, high expected return projects would remain in the applicant pool. If such a high-risk borrower is successful, the loan is repaid; if not, the lender simply gets stuck. This is adverse selection in the capital market. At some point, the gains from a higher interest rate will be offset by the losses because of the lower quality of the borrower pool, at which the (Nash) equilibrium interest rate will be determined. Note that there is an excess demand for credit at this interest rate.

Consumer Goods: In a traditional consumer good market, the price is set to equate supply and demand. But suppose the quality of the good is difficult to measure, and the higher the price, the higher the expected quality of the good. Then consumers will be willing to pay a higher than market clearing price for the good, trading off the higher price for the higher expected quality. Once again, the (Nash) equilibrium involves an excess supply of the good.

For an overview of adverse selection models and other failures of markets to clear because of private information, see Stiglitz (1987). Such models have been successfully applied to building a microeconomic foundation for some of the macroeconomic phenomena described by John Maynard Keynes (Mankiw and Romer 1991). The analytical theory of adverse selection was initiated by George Akerlof's (1970) famous "The Market for 'Lemons'." The next example is in the spirit, but not the letter of Akerlof's contribution.

12.5 A Market for Lemons

Consider a used car market, in which the sellers all know the value of the cars they want to sell, but the buyers only know the distribution of car values. Suppose all agents are risk neutral (§16.22) and the value v of a used car is a random variable uniformly distributed on the interval $[v_*, v^*]$, where $0 < v_* < v^*$, and the reservation price for the seller of a car of value v is fv, where $0 < f < 1$. For simplicity, we assume this is the same for all sellers, and all buyers know f.

A pure strategy for a buyer is thus the choice of a price offer p. Clearly it is a best response for a seller to accept any price $p \geq fv$. Assume there is perfect competition, so each buyers offers a price $p = \mathbf{E}[v|p]$; i.e., the price equals the expected value of the car, assuming the seller is willing to trade at that price.

For what values of f is this market Pareto-efficient? For what values of f are no cars bought or sold? What fraction of cars are sold in general?

To answer these questions, first we show that the allocation is Pareto-efficient if and only if

$$f \leq \frac{1 + v_*/v^*}{2}. \tag{12.2}$$

Proof: If the equilibrium is Pareto-optimal, then all cars must be sold, so $p \geq fv^*$ must hold. The expected value of a car in this case is $\mathbf{E}[v|p] = (v^* + v_*)/2$, so $p = (v^* + v_*)/2$. Thus $p = (v^* + v_*)/2 \geq fv^*$, or $f \leq (1 + v_*/v^*)/2$.

Now suppose $f \leq (1 + v_*/v^*)/2$ and let $p = (v^* + v_*)/2$. Then $fv^* \leq (v^* + v_*)/2 = p$, so $\mathbf{E}[v|p] = p$, and the equilibrium is Pareto-optimal. Q.E.D.

We also can show that if (12.2) fails, then the equilibrium price is given by

$$p = \frac{v_*}{2 - 1/f}.$$

The range of cars on the market will be $[v_*, v_*/(2f - 1)]$. In particular, as f approaches 1, this interval shrinks to zero length. This is Gresham's law in action: bad cars drive out good.

Proof: If (12.2) fails (i.e., if cars are relatively valuable to their owners), then it we can show that for a price $p \geq v_*$, the cars offered for sale will be uniformly distributed on $[v_*, p/f]$. To see this we use Bayes Rule. Let $V(\tilde{v})$ be the event "true value of car $\leq \tilde{v}$", and let $S(p)$ be the event "owner willing to sell car for price p." Then for $fv_* \leq p < fv^*$ we have

$$\begin{aligned}
\mathrm{P}[V(\tilde{v})|S(p)] &= \frac{\mathrm{P}[V(\tilde{v})\&S(p)]}{\mathrm{P}[S(p)]} \\
&= \frac{\mathrm{P}[S(p)|V(\tilde{v})]\mathrm{P}[V(\tilde{v})]}{\mathrm{P}[S(p)]} \\
&= \frac{\mathrm{P}[S(p)|V(\tilde{v})](\tilde{v} - v_*)/(v^* - v_*)}{(p/f - v_*)/(v^* - v_*)} \\
&= \frac{\mathrm{P}[S(p)|V(\tilde{v})](\tilde{v} - v_*)}{p/f - v_*}
\end{aligned}$$

which is zero for $\tilde{v} > p/f$ and $(\tilde{v} - v_*)/(p/f - v_*)$ for $fv_* \leq \tilde{v} \leq p/f$. This is the uniform distribution on $[v_*, p/f]$.

Prospective buyers in this case will purchase if and only if $p \leq (p/f + v_*)/2$, so the only possible Nash equilibrium has $p = (p/f + v_*)/2$, or

$$p = \frac{v_*}{2 - 1/f}.$$

The range of cars on the market will be $[v_*, v_*/(2f - 1)]$, as asserted.

12.6 Choosing an Exorcist

If your body is inhabited by an evil spirit, you go an exorcist. Suppose there are two types of exorcists: thorough and casual. Exorcists know their type. The possessed (or their loved ones) can costlessly inspect exorcists, and with probability $q = 10\%$ inspection reveals a thorough

exorcist to be thorough, while with probability $1-q$ the inspection of a thorough exorcist is inconclusive. Inspecting a casual exorcist is always inconclusive. Casual exorcism is worth \$100 to the buyer and costs the seller \$50 to provide, while thorough exorcism is worth \$600 to the buyer and costs the seller \$300 to provide. If an exorcist offers spells at price p, and if the buyer accepts, the buyer's payoff is the difference between the value of the spells and p, and the payoff to the exorcist is p minus the cost of conjuring the spells.

Suppose that the equilibrium price p^* of exorcism that is claimed to be thorough is determined by the condition that the return to all exorcists from an encounter with a possessed buyer be equal. Suppose also that buyers and sellers of exorcist spells meet randomly.

By eliminating dominated strategies, we need only consider only the following. Buyers accept any price $\leq$ \$100. A buyer who discovers (with probability q) that a seller is thorough always accepts p^*, and in addition, of buyers who do not know the type of a seller asking p^*, a fraction α accept p^*, and the remaining fraction $1-\alpha$ reject p^*. Thorough exorcists all offer to sell at p^*, but a fraction β of casual exorcists also offer to sell at p^*, misrepresenting their spells as thorough, while the remaining fraction $1-\beta$ of casual exorcists market their services truthfully. It so happens, and it is common knowledge, that 25% of exorcists are casual.

We can find the appropriate values of p^*, α and β such that the above represents a Nash equilibrium as follows. First, we represent the expanded game tree in Figure 12.5. Notice that we have pruned the game tree according to the characteristics of the three types of exorcists, in order to keep the diagram manageable.

We can find α by noting that the payoff to honest casual, dishonest casual, and thorough exorcists must be equal. The payoff to honesty is \$100 − \$50 = \$50, so the payoff to selling at p^* must also be \$50 for the dishonest exorcist. But this payoff is just $\alpha(p^*-\$50)$, so $\alpha = 50/(p^*-50)$. The payoff to a thorough exorcist is then $(q+(1-q)\alpha)(p^*-300)$, which must also equal 50. This give the quadratic equation $qp^2 - 400qp + 30000q - 12500 = 0$ for p^*, with solution

$$p^* = 200 + 100\sqrt{1+1.25/q}.$$

Now we have

$$\alpha = 50/(150 + 100\sqrt{1+1.25/q}).$$

With $q = 10\%$, this gives $p^* = 567.42$ and $\alpha = 9.66\%$.

To find β—the fraction of cheaters among the casual exorcists, we proceed as follows. Suppose inspection of a seller reveals nothing. Let

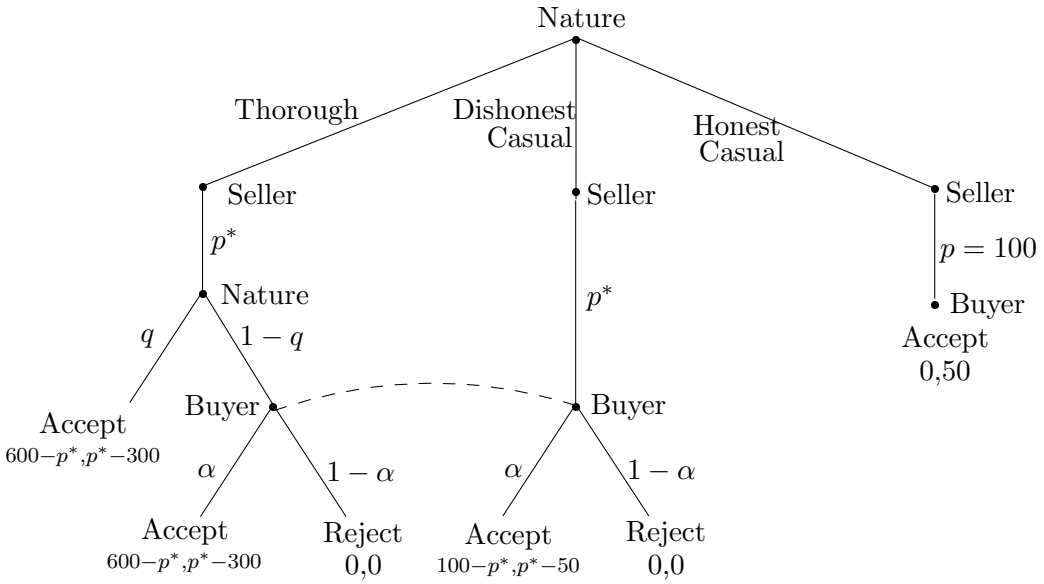

Figure 12.5: Choosing an Exorcist: Example of a Bayesian Game

$\pi(p)$ be the probability that the spells are casual given the asking price is p. Using Bayes' Rule, we find

$$\begin{aligned}\pi(p) &= \mathrm{P}[\text{casual}|p] \\ &= \frac{\mathrm{P}[p|\text{casual}]\mathrm{P}[\text{casual}]}{\mathrm{P}[p|\text{casual}]\mathrm{P}[\text{casual}] + \mathrm{P}[p|\text{thorough}]\mathrm{P}[\text{thorough}]}.\end{aligned}$$

Using the fact that 25% of exorcists are casual, we get

$$\pi(p^*) = \frac{\beta \cdot \frac{1}{4}}{\beta \cdot \frac{1}{4} + 1 \cdot \frac{3}{4}} = \frac{\beta}{\beta + 3}, \tag{12.3}$$

and of course $\pi(\$100) = 1$ (a exorcist selling services for \$100 is surely casual). Since only a fraction of buyers accept p^*, the payoff to accept and reject must be equal. The payoff to reject is zero, so the payoff to accept must also be zero. But the expected payoff to accept is

$$100\pi(p^*) + 600(1 - \pi(p^*)) - p^*,$$

so

$$\pi(p^*) = \frac{4}{5} - \frac{1}{5}\sqrt{1 + 1.25/q},$$

so $\pi(p^*) = 6.51\%$. Substituting this in (12.3), we get $\beta = 3\pi(p^*)/(1 - \pi(p^*)) = 20.90\%$: a casual exorcist is dishonest with probability 20.9%.

12.7 A First Price Sealed-Bid Auction

An object is to be sold to the highest bidder in a *first price sealed-bid auction* (i.e., bids are submitted simultaneously, the highest bidder wins, and pays the highest bid). Suppose there are two bidders, for each of whom the value v of the object is uniformly and independently distributed on the unit interval. This is thus a Bayesian game where the type of each player is the value of the object. Suppose both players choose a strategy of bidding $f(v)$, an increasing differentiable function of v. Given value v_1, Bidder 1 wants to choose b_1 to maximize

$$(v_1 - b_1)\mathrm{P}[b_1 > f(v_2)] = (v_1 - b_1)\mathrm{P}[g(b_1) > v_2] = (v_1 - b_1)g(b_1),$$

where $g(b) = f^{-1}(b)$, which exists since $f(v)$ is strictly increasing. Thus given v_1, b_1 satisfies the first order condition $g(b_1) = (v_1 - b_1)g'(b_1)$. By the symmetry of the problem, however, we must have $v_1 = g(b_1)$, so we have the differential equation

$$g(b_1) = (g(b_1) - b_1)g'(b_1).$$

This equation is satisfied by the linear function $g(b) = 2b$.[4] It follows that each bidder will bid $b = v/2$.

Now suppose there are k bidders, each with an independently distributed value v on the unit interval. Bidder 1 chooses b_1 to maximize

$$\begin{aligned}(v_1 - b_1)\mathrm{P}[b_1 &< \min\{f(v_2), \ldots, f(v_n)\}] \\ &= (v_1 - b_1)\mathrm{P}[g(b_1) < \min\{v_2, \ldots, v_n\}] \\ &= (v_1 - b_1)g(b_1)^{k-1}.\end{aligned}$$

The first order condition, after a bit of simplification, is now

$$g(b_1) = (k-1)(g(b_1) - b_1)g'(b_1).$$

This is satisfied by $g(b) = kb/(k-1)$, so the optimal bid is $f(v) = (k-1)v/k$. Note that as the number of bidders increases, the optimal bid approaches v.

[4]To arrive at this equation, assume $g(b)$ is a polynomial of some low degree and equate the coefficients of b on both sides of the differential equation to show that $2b$ works. I don't know if this solution is unique, though.

12.8 A Common Value Auction: The Winner's Curse

In the second price auction (§2.8), the increasing bid auction (§2.10), and the first price sealed-bid auction (§12.6), each bidder has a private value for the object under contention, and this value is known with certainty. The *common value auction* is a first-price, sealed bid auction in which the object has the same value for all bidders, but they do not know that value—for instance, bidding for oil drilling rights.

It is easy to see that in a common value auction, bidding your *measured value* will get you into trouble, since if you win, you are more likely than any other bidder to have overestimated the value! For instance, suppose there are n bidders, each of whom is allowed to take one 'measurement' of the value, and suppose the measured value is $\tilde{v} = v + \epsilon$, where v is the actual value and ϵ is an error term that is uniformly distributed on the closed interval $[-a, a]$. Suppose all bidders bid their measured values, and suppose you win. Prove that your expected loss is

$$a\frac{n-1}{n(n+1)} > 0, \tag{12.4}$$

so you lose by winning! This is called the *winner's curse*. Note that while the expected loss per bidder fall to zero as the number n of bidders increases, the total loss of all bidders is approximately a.

Since professional bidder are not usually stupid, they must bid less than their measured value in a common value auction. One possibility it to bid a constant amount less. But it is easy to see that this does not work. For instance, suppose everyone but you bids an amount b less than the measured value, and you bid c less than your measured value, choosing c to maximize your expected profit. You can check that your profit is proportional to

$$-(2a+b-c)^n(a(n-1)-b+c) - (b-c)^n(a(n+1)+b-c).$$

Since all bidders make the same calculation, we must have $b = c$, in which case this expression becomes $-2^n a^{n+1}(n-1)$, which is negative. So this does not solve the problem of how to bid. You might check the literature to see what solutions are proposed to the Winner's Curse.

12.9 Common Value Auction II

The government ordered a quantum spin decoder built for its SETI (Search for Extra Terrestrial Intelligence) project, and put it up for

auction when the project was abandoned. The decoder is commercially useless, but it has two potentially valuable parts, a kramel inhibitor and a beta-phase detector. General Electric and Martin Marietta sent technicians to check out the decoder prior to the auction, but the GE technician didn't know much about beta-phase detectors, and reported back that the kramel inhibitor is worth x on the market, and the beta-phase detector is worth at most $\beta > 0$. The Martin Marietta technician didn't know much about kramel inhibitors, and reported back that beta-phase detector is worth y on the market, and the kramel inhibitor is worth at most $\alpha > 0$.

So suppose the value of the decoder is $\tilde{x} + \tilde{y}$ where $\tilde{x}$ is uniformly distributed on $[0, \alpha]$ and $\tilde{y}$ is uniformly distributed on $[0, \beta]$. Suppose GE knows the precise value x of the kramel inhibitor, and Martin Marietta knows the precise value y of the beta-phase detector. Suppose, moreover, that these facts are common knowledge to both firms.

(a) Suppose Martin Marietta sends an industrial spy to the economics division of GE, and returns triumphantly with Martin Marietta's bidding formula:

$$\left(\frac{\alpha+\beta}{2}\right)\frac{x}{\alpha}. \tag{12.5}$$

"But what is x?" asks her boss. "That I could not find out," she replies. Similarly, GE sends out a spy who returns with the formula

$$\left(\frac{\alpha+\beta}{2}\right)\frac{y}{\beta} \tag{12.6}$$

but does not know y. Show that these strategies form a linear Bayesian Nash equilibrium, find the value of the auction to the two firms, and find the probability that each wins.

(b) How do you make sense of the two strategies?

(c) It looks like we are getting somewhere with this problem, but we're really not. To see this, suppose GE's spy came back and said, "Well, you know how our economists came up with the number $\delta > 1$ such that we should bid δx? Well, I found out Martin Marietta also has a magic number γ, that they use. They will bid γy, and moreover, you can check that $\gamma = \delta/(2\delta - 1)$." Strange but true. The Martin Marietta spy comes back with a similar story of having found out GE's δ, and $\delta = \gamma/(2\gamma - 1)$.

Show that these two strategies form a linear Bayesian Nash equilibrium, whatever $\gamma > 1$ is.

(d) Things are a bit better if there are three bidders. Show that in this case the strategy of each firm bidding 5/3 times its observation is a linear Bayesian Nash equilibrium.

(e) Show that if there are four bidders, there is a unique linear Bayesian Nash equilibrium where each firm bids 9/4 times its observation. Only do this if you have a computer algebra program or a very sharp pencil

(f) Show that if there are five bidders, there is a linear Bayesian Nash equilibrium where each firm bids 14/5 times its observation. Only do this if you have a sophisticated computer algebra program or a very sharp pencil. I could not prove that this is unique, though I didn't try very hard.

12.10 Predatory Pricing: Pooling and Separating Equilibria

There are two firms, the 'incumbent' and the 'entrant'. There are also two periods. The incumbent had a monopoly in the industry previous to period 1, earning monopoly profits m in each period. But the entrant decided to enter the market in period 1. The incumbent takes an action during period 1, either 'prey' (charge a very low price to drive the entrant out of the market) or 'accommodate' (charge a profit-maximizing price). The entrant's profit is $d_2 > 0$ if the incumbent accommodates and his loss is $p_2 > 0$ if the incumbent preys, with $m > d_2$. The incumbent may be either 'sane' or 'crazy.' When sane, the incumbent makes d_1 by accommodating and loses p_1 when preying, with $d_1, p_1 > 0$. If the incumbent is crazy, he loses money by preying, but reaps the psychic reward of punishing the entrant for disturbing his monopoly (or perhaps he has a different assessment of the costs and benefits of maintaining the monopoly, or has a reputation to uphold). Thus if the incumbent is crazy he always preys.

However even a sane incumbent might decide to prey, to try to convince the entrant that the incumbent is crazy and drive him out of the market in the second period. Let π_1 be the entrant's prior probability that the incumbent is sane (this is a Bayesian subjective probability, perhaps based on the frequency of sane incumbents in the economy), before the incumbent takes an action, and let π_2 be the entrant's posterior judgement that the incumbent is sane if he preyed in the first period (we assume that an incumbent is sane with probability 1 if he accommodated in the first period).

In period 2, the entrant chooses an action, either 'stay' or 'exit.' If the incumbent is sane, he will not prey in period 2, since it can have no effect on the entrant's behavior. So by staying the entrant will earn d_2 if the incumbent is sane, and lose p_2 if the incumbent is crazy. If the incumbent is sane, he earns d_1 in period 2 if the entrant stays, and $m > d_1$ if the entrant exits.

Figure 12.6 shows a partial game tree for the problem, showing only the part of the game tree where the incumbent is sane, and assuming that the discount factor for second period payoffs is δ.

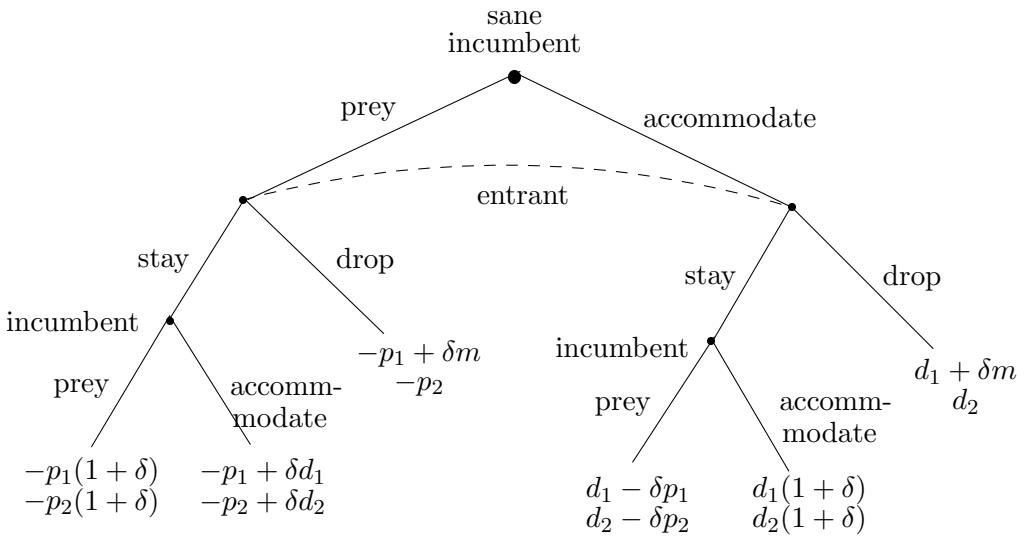

Figure 12.6: The Predatory Pricing Model: This shows the part of the game tree starting from the Sane Incumbent. Note that all the nodes for the Entrant are part of two-node information sets, the other node of which is 'Crazy Incumbent,' who always preys.

We say a perfect Bayesian equilibrium is *separating* if the sane firm accommodates in the first period. Otherwise the equilibrium is *pooling*. A separating equilibrium is one in which the 'type' of a player (e.g., sane/crazy) is revealed by the player's behavior, whereas a pooling equilibrium is one in which different types of players behave in the same way. In a separating equilibrium for this game, the entrant has perfect knowledge in the second period.

(a) Show that the condition for a separating equilibrium is

$$-p_1 + \delta m \leq d_1(1+\delta). \tag{12.7}$$

(b) Show that the condition for a pooling equilibrium is

$$\pi_1 \leq \frac{p_2}{d_2 + p_2} \tag{12.8}$$

(c) If neither of these conditions holds, show that we may have a 'hybrid' perfect Bayesian equilibrium, in which incumbent and entrant both randomize their choices. Show that this mixed strategy equilibrium is given as follows. Let α be the probability that a sane incumbent preys in the first period, and let β be the probability that the entrant stays in the second period if the incumbent preys in the first period. Then we have

$$\alpha = \frac{(1 - \pi_1)p_2}{\pi_1 d_2}, \tag{12.9}$$

$$\beta = \frac{-p_1 + \delta m - d_1(1 + \delta)}{m - d_1}. \tag{12.10}$$

12.11 Limit Pricing

Pre-game theoretic industrial organization theorists often held that a monopolist would use *anticipatory pricing*, maintaining a product price lower than the monopoly price, in order to deter competition from entering the market. In anti-trust law, it was hard to apply that argument against a monopolist, because it suggested that the monopolist be forced to raise his price!

But why bother lowering your price *before* an entrant comes into the market? Perhaps predatory pricing is a good idea, but it is certainly not clear that anticipatory pricing is. But suppose the incumbent wants to signal to a potential entrant that his costs of production are so low that the entrant couldn't make a profit by entering the market. Could this be a reasonable strategy? Let's see.

Suppose there is an incumbent firm with monopoly power in period 1 and this firm chooses first period price p. Then in period 2, firm 2 decides whether to enter or to stay out. If firm 2 enters, there is a duopoly for that period, and the game is over. Otherwise firm 1 remains a monopoly.

Suppose both firms have constant marginal costs and no fixed costs. Moreover marginal cost takes two values, high ($c = c_h$) or low ($c = c_l$). The entrant has high costs, and the incumbent can have low costs with probability q, or high costs with probability $1 - q$. Let $m^h(p)$ be the incumbent's monopoly profit when he charges p and $c = c_h$, and let $m^l(p)$

be the incumbent's monopoly profit when he charges p and $c = c_l$. Let p_m^l and p_m^h denote the monopoly prices charged by the incumbent when marginal cost is low or high (i.e., the monopolist's profit-maximizing price).

Let $m^l = m^l(p_m^l)$ and $m^h = m^h(p_m^h)$ be the incumbent's profit when he maximizes short-run profits, depending on his type. Suppose the incumbent knows his cost, and the entrant learns the incumbent's cost after entering the market and committing himself to production. Thus the duopoly price is independent of first-period prices. We denote the entrant's duopoly profits by d^h when costs are high and d^l when costs are low. We assume that $d^h > 0 > d^l$, so the entrant would enter only if the incumbent's costs are high (otherwise the problem is trivial).

The only way the incumbent can signal low cost is by charging a low price p^l. If $c = c_l$ and $p^l = p_m^l$, this is clearly optimal. But even if $c = c_h$, by charging p^l, the monopolist might convince the entrant not to enter, so the monopolist could charge the monopoly price in the second period. Thus the l-firm might charge $p^l < p_m^l$ to deter the h-firm from entering. This is know as *limit pricing*.

Of course, if the entrant believes that a monopolist would charge a low price even if $c = c_h$, then observing the low price would have no effect on the entrant's decision. As before, we say a perfect Bayesian equilibrium is a *separating equilibrium* if the incumbent's action reveal his type; otherwise we say the equilibrium is a *pooling equilibrium*.

(a) Prove that $p_m^l < p_m^h$, so it is plausible that a monopolist charging a low price has low costs.

(b) Under what conditions do we have a separating equilibrium?

(c) Give an example of a separating equilibrium.

We conclude from this analysis that when there is a separating equilibrium (a) because there is limit pricing (the incumbent manipulates his price to prevent pooling), the entrant knows the type of the incumbent perfectly; and (b) consumers are better off than if the incumbent were prevented from using limit pricing, since they get a lower price in the first period, and entry still occurs when it should (i.e., when the incumbent is a high-cost producer).

12.12 A Simple Limit Pricing Model

Consider a two-period game in which in the first period firm A has a monopoly position in an industry in which demand is $q = m - p$ where p is price and q is quantity. Suppose the firm has constant marginal cost c, which can either be high (c_h) or low ($c_l < c_h$). Firm A maximizes profits $\pi = (p - c)q = (m - q - c)q$ by setting $q = (m - c)/2$, and $p = m - q = (m + c)/2$.

Suppose however there is a firm B which has constant marginal cost c_h and fixed cost f, and which could enter the market in the second period. Should B enter, the Cournot duopoly price will obtain in the second period. We assume that A knows B's production costs, but B does not know whether A is a high cost (c_h) or a low cost (c_l) producer. If A charges the monopoly price in the first period, then of course B can use this price to infer whether or not A is a low or high cost producer. To make the problem interesting, we assume that B would have negative profits competing against A if A is a low cost producer, and positive profits if A is a high cost producer. But of course A could charge a price different from his monopoly price in the first period to deter B from entering in the second period.

Under what conditions will there be a separating equilibrium, in which A charges his monopoly price in the first period? Note: you can leave your solution in the form of inequalities that must be satisfied in a separating equilibrium (be sure to define separating and pooling equilibria, including the belief assumptions involved).